Sensitivity Analysis of Flexible Provisioning

Sebastian Stein, Terry R. Payne and Nicholas R. Jennings

IAM Research Group,
School of Electronics and Computer Science,
University of Southampton,
Southampton, SO17 1BJ, UK.

E-mail: {ss04r,trp,nrj}@ecs.soton.ac.uk

Technical Report Number: ECSTR-IAM07-001
ISBN: 978 085432 870 3

Abstract

This technical report contains a sensitivity analysis to extend our previous work. We show that our flexible service provisioning strategy is robust to inaccurate performance information (when the available information is within 10% of the true value), and that it degrades gracefully as the information becomes less accurate. We also identify and discuss one particular case where inaccurate information may lead to undesirable losses in highly unreliable environments.

Keywords: *Service-Oriented Computing, Service Provisioning, Workflows*

1 Introduction

Semantic Web services promise to revolutionise the way computational resources and business processes are offered and invoked in open, distributed systems, such as the Internet. These services are described using machine-readable meta-data, which enables consumer applications to automatically discover and provision suitable services for their workflows at run-time. However, current approaches have typically assumed service descriptions are accurate and deterministic, and so have neglected to account for the fact that services in these open systems are inherently unreliable and uncertain. Specifically, network failures, software bugs and competition for services may regularly lead to execution delays or even service failures.

To address this problem, we have proposed a flexible approach to service provisioning, which varies the provisioning of tasks that are part of complex workflows according to the performance characteristics of the service providers [1; 2]. In our previous work, we showed that our strategy achieves an almost four-fold improvement over static strategies that do not provision services in a flexible manner. However, so far, we have assumed accurate information to be available to the service-consuming agent (more specifically, we assumed accurate failure probabilities and service duration distributions). In this report, we conduct a sensitivity analysis to examine how our strategy performs in the presence of inaccurate information (e.g., when the service consumer has to base its decisions on limited previous interactions or even incorrect trust sources).

2 Experimental Setup

We follow the same experimental setup as described in Section 5.1 of [2], but now systematically introduce errors into the information that is available to the service consumer. More specifically, we define the following strategy:

Definition 2.1 (Inaccurate($\tilde{\lambda}, \tilde{\mu}$) Strategy)**.** The *inaccurate*($\tilde{\lambda}, \tilde{\mu}$) strategy provisions services as the *flexible* strategy, but relies on wrong information about the performance of service providers: instead of the real failure probability f_i, it uses $f'_i = \min(1, \tilde{\lambda} f_i)$, and instead of $d_i = \text{Gamma}(k, \theta)$, it uses $d'_i = \text{Gamma}(k, \tilde{\mu}\theta)$ for provisioning[1].

Hence, an agent following the *inaccurate(0.8,2)* strategy would *underestimate* the failure probability of providers as 80% of the actual value, and *overestimate* the duration (the perceived mean is doubled). The *inaccurate(1,1)* strategy is equivalent to the *flexible* strategy.

In the following section, we present the results of a number of experiments with varying values for $\tilde{\lambda}$ and $\tilde{\mu}$ to investigate the sensitivity of our strategy. Generally, we expect the performance of the *inaccurate*($\tilde{\lambda}, \tilde{\mu}$) strategy to decrease as their perceived performance information becomes less accurate (i.e., as $\tilde{\lambda}$ or $\tilde{\mu}$ diverge further from 1).

3 Results

Figure 1 shows the performance of our strategy when it underestimates the failure probability of service providers[2]. In most cases, the average utility

[1] We assume that a Gamma distribution is used for service durations.

[2] As in our previous work, we use the average net profit of a service consumer to evaluate its performance. This is the difference of the reward gained from completing the workflow (if any) and the costs incurred.

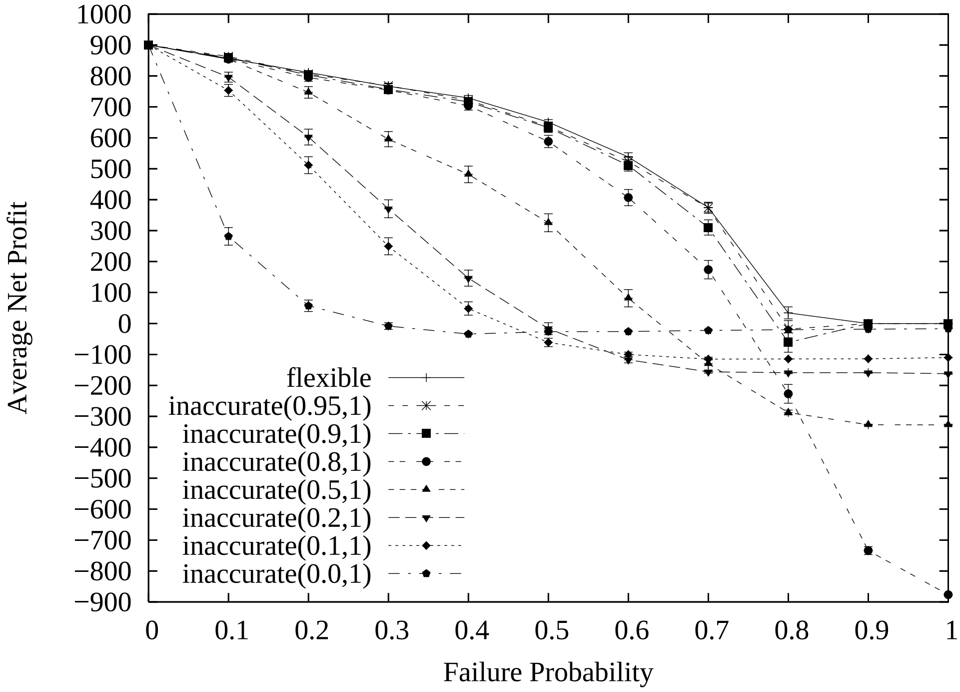

FIGURE 1: Average net profit when underestimating the failure probability of providers ($\tilde{\lambda} < 1$).

gained by the strategy degrades gracefully as the performance information becomes more inaccurate. In fact, when the (true) failure probability is low in the environment (up to around 0.3), even the highly inaccurate strategies ($\tilde{\lambda} = 0.2$ and $\tilde{\lambda} = 0.1$) do well. However, when the failure probability rises to 0.8 and beyond, some of the inaccurate strategies perform very poorly. This includes *inaccurate(0.8,1)*, which achieves a very low utility at high failure probabilities. This is because it provisions a large number of providers in parallel without detecting that the workflow is infeasible (and thus, it usually loses its high investment). Despite this, the results also show that small inaccuracies in the information have little or no affect on our strategy (up to around 10 %).

Figure 2 shows the corresponding data when our strategy underestimates the duration of service providers. In these settings, our strategy handles an error of up to 20%-30% ($\tilde{\mu} \leq 0.7$) very well — the results are only marginally worse than the accurate *flexible* strategy. Beyond that, the performance drops significantly and sometimes erratically. This is not surprising, because the inaccurate strategies will use low waiting times, and so increasingly invoke new services before waiting for the previous ones to finish successfully.

In further experiments, we evaluated the effect of overestimating various performance parameters. Figure 3 shows the results[3] of the *inaccurate* strategy when

[3] Note that the confidence intervals are generally smaller in these cases. This is due to unnecessary overprovisioning, resulting in a smaller overall profit variance.

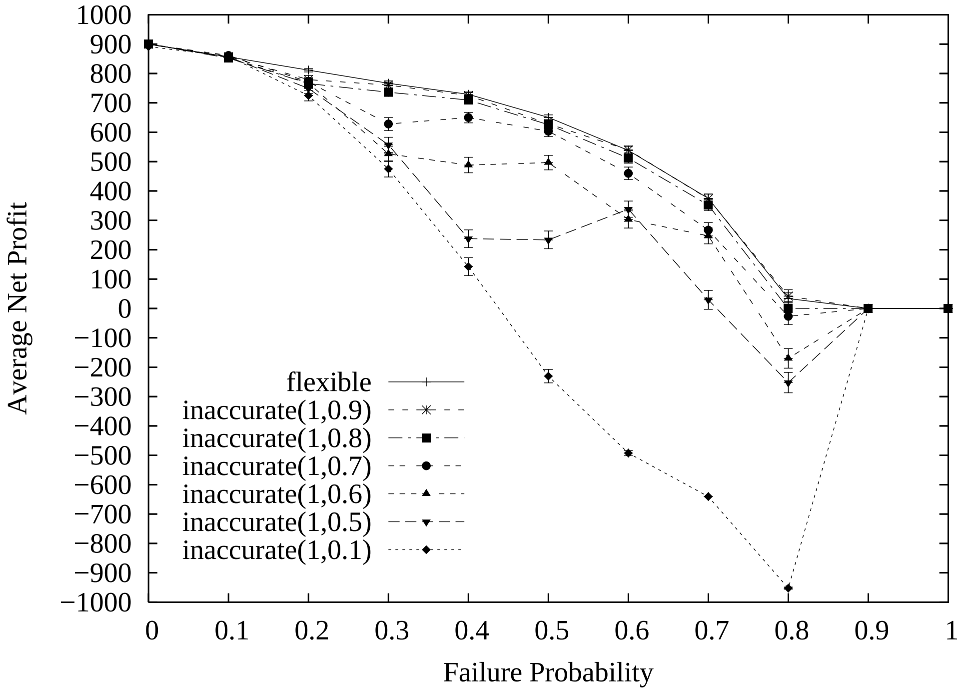

FIGURE 2: Average net profit when underestimating the service duration of providers ($\tilde{\mu} < 1$).

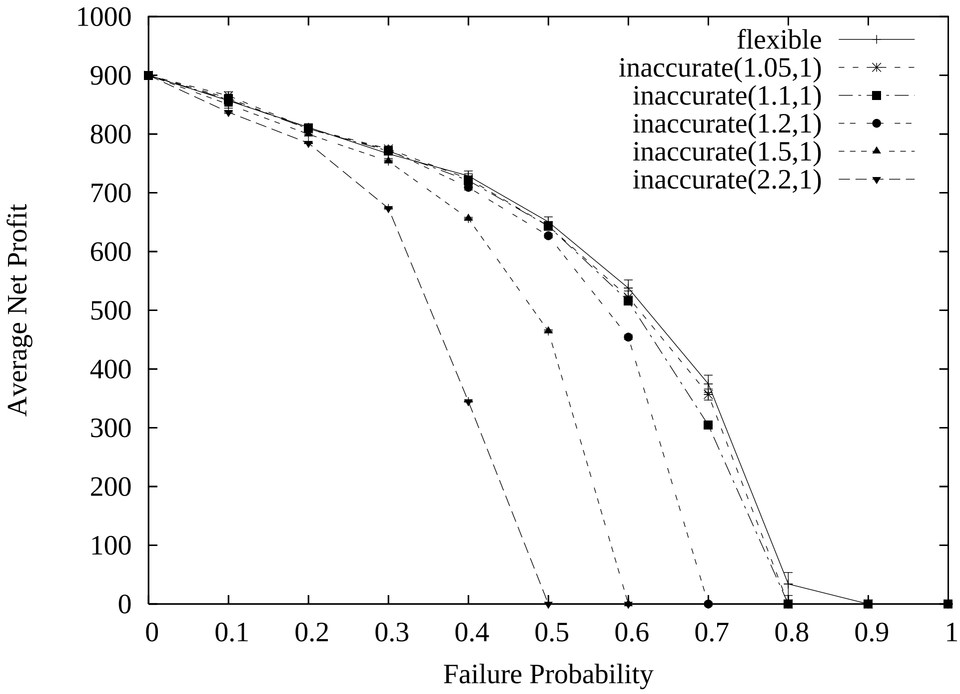

FIGURE 3: Average net profit when overestimating the failure probability of providers ($\tilde{\lambda} > 1$).

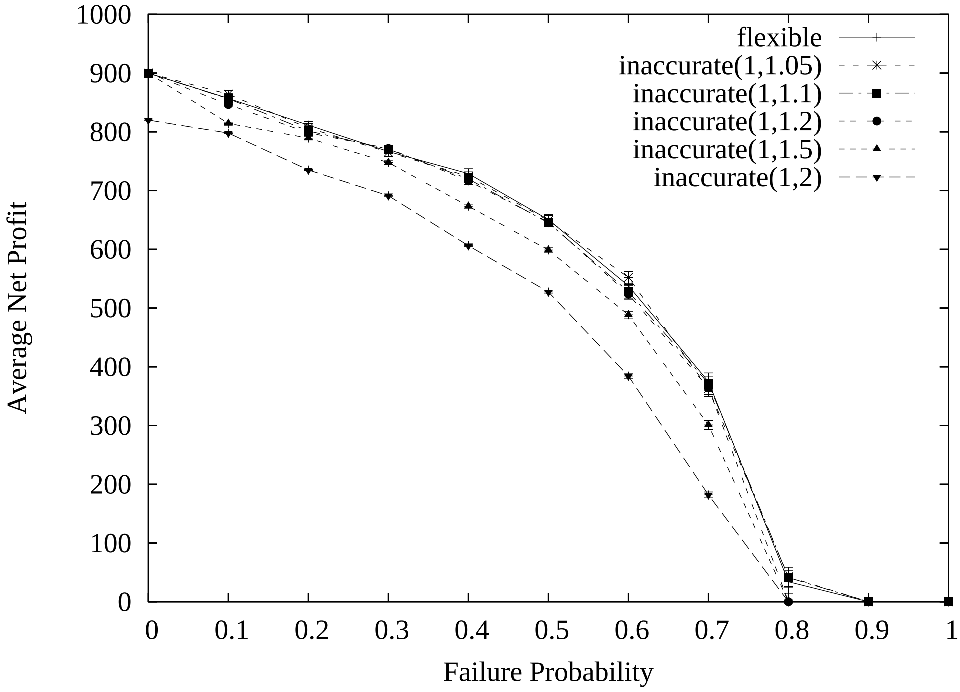

FIGURE 4: Average net profit when overestimating the service duration of providers ($\tilde{\mu} > 1$).

$\tilde{\lambda} > 1$. Not surprisingly, the performance of the strategy simply degrades as the perceived failure probability rises. Because it is inherently more conservative (it will provision unnecessarily many providers), it never makes a long-term loss. Figure 4 shows the corresponding results when the consumer overestimates the service duration. Here, the performance again degrades slowly as the duration rises. This is because the agent allocates unnecessarily long timeouts to services. However, the loss in performance is not as pronounced as in the previous example, because the agent will simply continue in its workflow when a service succeeds earlier than expected (and thus, no resource is wasted).

4 Conclusions

The results presented in this report show that our strategy is robust to small and moderate inaccuracies. In all cases, it performs well when the information provided is within 10% of the true value. Otherwise, performance usually degrades gracefully as larger errors are introduced into the information that is known about providers (until they are too large to be of any use to the agent — e.g., as $\tilde{\mu}$ reaches 0.5). A notable exception is the case when the overall failure probability in the system is very large. Here, inaccuracies may lead the agent to highly over-provision tasks in order to achieve a small expected profit. Wrong information about service failures then means that the agent will spend a large amount of resources on invoking services, but still fail to complete the workflow. This result indicates that our strategy may benefit from abandoning highly over-provisioned workflows in environments where information is inaccurate, and we will consider this in future work. Otherwise, the results presented here are promising and show that our strategy is applicable even in environments where completely accurate performance information is unavailable (as will be typical in a large multi-agent system, where such information may come from third-party trust and reputation systems).

References

[1] S. Stein, N. R. Jennings, and T. R. Payne. Flexible provisioning of service workflows. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI-06), Riva del Garda, Italy*, pages 295–299. IOS Press, 2006.

[2] S. Stein, T. R. Payne, and N. R. Jennings. Flexible provisioning of semantic web service workflows. In *ACM Transactions on Internet Technology*, 2007. (submitted).